

A note on hierarchical hubbing for a generalization of the VPN problem

Neil Olver*

July 1, 2014

Abstract

Robust network design refers to a class of optimization problems that occur when designing networks to efficiently handle variable demands. The notion of “hierarchical hubbing” was introduced (in the narrow context of a specific robust network design question), by Olver and Shepherd [14]. Hierarchical hubbing allows for routings with a multiplicity of “hubs” which are connected to the terminals and to each other in a treelike fashion. Recently, Fréchet et al. [8] explored this notion much more generally, focusing on its applicability to an extension of the well-studied hose model that allows for upper bounds on individual point-to-point demands. In this paper, we consider hierarchical hubbing in the context of a different (and extremely natural) generalization of the hose model, studied earlier in [14], and prove that the optimal hierarchical hubbing solution can be found efficiently. This result is relevant to a recently proposed generalization of the “VPN Conjecture”.

1 Introduction

1.1 Robust network design

Robust network design considers the problem of building networks under uncertainty in the pattern of utilization. Introduced by Ben-Ameur and Kerivin [1], the framework encompasses the important case of the “hose model” introduced by Fingerhut [6] and Duffield et al. [4]. It can itself be seen as falling under the broader umbrella of robust optimization [2].

*Department of Econometrics & Operations Research, VU University Amsterdam, and Centrum Wiskunde & Informatica (CWI), Netherlands. Email: n.olver@vu.nl. Research supported by a NWO Veni grant.

We refer the reader to [13] for a more in-depth treatment; here we will give a brief self-contained exposition of the model. We are given an undirected graph $G = (V, E)$; this should be interpreted as an existing high-capacity network, in which we can reserve capacity. We assume there is an unlimited total capacity on any given link of the network, and that the cost to reserve capacity on any link is a linear function of the capacity required. Let $c : E \rightarrow \mathbb{R}_+$ denote the per-unit cost of capacity on each edge. A set $W \subseteq V$ of *terminals* need to be adequately connected using the capacity reserved.

A *traffic pattern* (or *demand pattern*) describes the precise pairwise demand requirements at some moment in time. It can be specified by a traffic matrix D , indexed by pairs of terminals; for terminals i, j , the entry D_{ij} represents the bandwidth needed to send data from i to j . In our network, the traffic pattern is not fixed, but varying (and possibly uncertain). To deal with this, the robust network design framework allows for a *set* of traffic patterns to be prescribed. This (it turns out) can always be taken to be a convex set, and so we describe this set, or *demand universe*, as a convex body $\mathcal{U} \subset \mathbb{R}_+^{W \times W}$.

The robust network design (RND) problem asks for the cheapest capacity reservation $u : E \rightarrow \mathbb{R}_+$ that can support all traffic patterns in the specified universe \mathcal{U} . To fully specify the problem however, a further aspect must be considered: the *routing scheme*. The coarsest division is into *oblivious* or *dynamic* routing. In dynamic routing, the way in which traffic is routed may vary arbitrarily as a function of the current traffic pattern. This is typically infeasible, and we will be concerned here with the more practical oblivious routing, where the routing used for any given pair of terminals is specified in advance. We will also only consider *single-path routing*. The routing scheme in this case is described by a template $\mathcal{P} = \{P_{ij} : i, j \in W\}$, where P_{ij} is an i - j -path for each $i, j \in W$. (We do not require this path to be simple.)

We may summarize the general robust network design problem (with oblivious, single-path routing) as follows:

RND problem. Given an undirected graph $G = (V, E)$ with edge costs $c(e)$, a terminal set $W \subseteq V$, and a convex demand universe $\mathcal{U} \subset \mathbb{R}_+^{W \times W}$, a solution to the robust network design problem consists of a routing template $\mathcal{P} = \{P_{ij} : i, j \in W\}$, and a capacity allocation $u : E \rightarrow \mathbb{R}_+$, such that \mathcal{U} can be routed according to \mathcal{P} within the capacity u , i.e.,

$$u(e) \geq \max_{D \in \mathcal{U}} \sum_{i,j \in W} D_{ij} \ell(P_{ij}, e) \quad \forall e \in E. \quad (1)$$

Here, $\ell(P, e)$ gives the number of times that edge e occurs on the (possibly non-simple) path P .

The difficulty in this optimization problem lies in choosing the routing template; once this is fixed, the optimal capacity allocation can be determined by solving a convex program described by (1), assuming we have access to at least a separation oracle for \mathcal{U} .

Note that there is always an optimum solution template whose paths P_{ij} are all simple, since any non-simple path can simply be replaced by a simple path within its support. The reason we allow non-simple paths is related to the specific type of routing templates we will be interested in.

A case of interest is that of *symmetric* demands; this means that demand from i to j is not distinguishable from demand from j to i . In this case, which will concern us in this paper, it is convenient to consider \mathcal{U} to be a subset of $\binom{W}{2}$, the set of unordered pairs of terminals, so that $D_{ij} = D_{ji}$ refers to the same demand, and $P_{ij} = P_{ji}$ to the same path. Equation (1) then becomes

$$u(e) \geq \max_{D \in \mathcal{U}} \sum_{\{i,j\} \subseteq W} D_{ij} \ell(P_{ij}, e). \quad (2)$$

The well-studied *symmetric hose model* [6, 4] is parameterized by a vector $b \in \mathbb{R}_+^W$, yielding the universe

$$\mathcal{H}(b) = \left\{ D \in \mathbb{R}_+^{\binom{W}{2}} : \sum_{\{i,j\} \subseteq W} D_{ij} \leq b_i \quad \forall i \in W \right\}.$$

This models the situation where terminals are connected to the network with “hoses” of known, fixed capacity, so that the total demand involving terminal i cannot exceed the capacity b_i of its associated hose link. Any demand pattern that fits through the hoses should be routable in the final network. These hoses may model real links, or chosen based on operational criteria; either way, the hose model gives a simple, useful and concise description of what the network must be able to handle, making it a very popular model in the literature.

A number of variations and generalizations of the symmetric hose model have been considered [5, 8, 14, 6]. For example, Fr chet te et al. [8] consider the “capped” hose model, where in addition to the hose capacities b , point-to-point upper bounds $\Gamma \in \mathbb{R}_+^{\binom{W}{2}}$ are also given, and the universe is

$$\mathcal{H}(b) \cap \left\{ D \in \mathbb{R}_+^{\binom{W}{2}} : D_{ij} \leq \Gamma_{ij} \quad \forall i, j \in W \right\}.$$

1.2 The generalized VPN problem

Rather than the capped hose model, we will be concerned with a different generalization of the hose model, introduced by Olver and Shepherd [14]. Let T^b be an arbitrary capacitated tree, with nonnegative edge capacities b and with leaf set in exact correspondence with the terminal set W . We will use T^b to define a demand universe in a simple and natural way: let $\mathcal{U}(T^b)$ consist of all demand patterns that can be routed on T^b .

The case where T^b is a star corresponds precisely to the hose model; the capacity of the edge adjacent to terminal i precisely gives the marginal of i . This generalization allows the network operator more precise control over the demand universe, hopefully leading to more efficient solutions. In particular, if the terminals of the network can be logically divided into distinct groups (e.g., different branches of the company), with limited communication between groups, this information can be encoded via $\mathcal{U}(T^b)$.

It was shown in [14] that the RND problem with oblivious routing for this class of demand universes is approximable to within a factor of 8. The algorithm that achieves this always returns a solution of a particular form—a *hierarchical hubbing*.

1.3 Hierarchical hubbing

Olver and Shepherd [14] proposed the following algorithm for the RND problem with universe $\mathcal{U}(T^b)$: find the cheapest embedding of T^b into the network. This embedding maps internal nodes of T^b into the network (these are “hubs”), and each edge e of T^b is mapped to a “cable” of capacity $b(e)$ that connects the hubs corresponding to the endpoints of the edge (see Figure 1). More than one node of T^b can be mapped to the same hub, and multiple cables may run over the same edge of the network. The routing template associated with a hierarchical hubbing is obtained, for each $\{i, j\} \subseteq W$, as the image of the unique i - j -path in the tree under the mapping (again, see Figure 1), yielding a (possibly non-simple) i - j -path in G . For any edge $e \in E(G)$, define the capacity $u(e)$ to be the sum of the cable capacities that use edge e . It is easy to show that this template and capacity allocation provides a valid solution to the RND problem with universe $\mathcal{U}(T^b)$.

What is very pleasant about this restricted class of solutions is that the cheapest such solution can be found in polynomial time. (The optimization problem is closely related to the zero-extension problem on trees [11, 3].) It is shown in [14] that the resulting solution is within a constant factor of the optimal solution to the RND problem for $\mathcal{U}(T^b)$, where arbitrary routing




Figure 1: An example of an embedding of T , with the resulting routing from 1 to 4 indicated.

templates are allowed.

Fr chet te et al. [8] generalize this notion in two ways. Firstly, they consider hierarchical hubbing templates for arbitrary demand universes, not just those associated with the generalized VPN problem. Secondly, they consider the optimization problem where the tree used to define the hierarchical hubbing is not specified, but may be chosen as part of the solution. More precisely, they define the following optimization problem:

RND_{HH} problem. Given an undirected graph $G = (V, E)$ with edge costs $c(e)$, a terminal set $W \subseteq V$, and a convex demand universe $\mathcal{U} \subset \mathbb{R}_+^{W \times W}$, a solution to the RND_{HH} problem consists of (i) a capacitated tree T^b , with leaf set W , such that any demand in \mathcal{U} can be routed on T^b , and (ii) an embedding of T^b . The capacity allocation $u(e)$ for an edge $e \in E$ is the sum of the capacities $b(f)$ of the cables that use edge e , and the goal is to minimize the final cost $\sum_{e \in E} c(e)u(e)$ of the solution.

It is easy to confirm that any solution to the RND_{HH} problem is a solution to the RND problem, but not vice versa. So in general the optimal solution to RND_{HH} can be more expensive than the optimal RND solution; in fact, Fr chet te et al. [8] demonstrate that the gap can be $\Omega(\log |V|)$, for some choices of the universe.

Fr chet te et al. [8] are motivated to consider hierarchical hubbing for a few reasons. In *hub routing*, all traffic is routed via a single hub node; this has the advantage that routing decisions are localized at the hub. In order to address some practical shortfalls of hub routing, Shepherd and Winzer [15] ask for a “multihub” extension of this. Fr chet te et al. argue that hierarchical hubbing provides a natural such extension (note that it is clearly a generalization; hub routing corresponds to taking the hub tree to be a star). They also show that it provides an effective heuristic for

finding good solutions to the capped hose model mentioned earlier, which is APX-hard in the single-path oblivious routing model. They observe that for the capped hose model (and for other universes as well), hierarchical hubbing can yield much cheaper solutions than using a single hub. This is in contrast to the vanilla hose model, where a hub routing provides an optimal oblivious routing solution (see Section 3).

Finding the optimal hierarchical hubbing solution for a given universe is in general APX-hard, as observed in [7] (cf. [12] for tree routings); this can be seen by choosing the universe so that the resulting RND problem is precisely the Steiner tree problem. It is thus natural to ask for which universes the problem is polynomially solvable. In fact, as essentially shown in [14] (see Lemma 9 later), the hierarchical hubbing problem is solvable exactly in polynomial time if the choice of hub tree is specified. The difficulty is thus in identifying the correct optimal choice of the hub tree. One case where the problem was previously known to be polynomial was for the hose model [10]; this will be discussed more in Section 3.

Remark 1. Any tree routing—meaning a routing template $\mathcal{P} = \{P_{ij} : i, j \in W\}$ such that $\bigcup_{i,j \in W} P_{ij}$ is a tree—can be described as a hierarchical hubbing. The hub tree is obtained from the support, adding additional edges as needed to ensure that all terminals are leaves in the hub tree.

Remark 2. It would also be natural to instead choose capacities by considering the routing template induced by the hierarchical hubbing, and using (2). This alternative formulation is in general not the same as described above; there may be situations where not all cables on a given edge can be simultaneously saturated by a traffic pattern in \mathcal{U} , leading to a larger capacity requirement with the cable formulation. The formulation that we use in this paper, and which is also used in [8], seems overall easier to deal with (see, e.g., Lemma 9 later). If the Generalized VPN Conjecture discussed in Section 3 is true, it follows immediately that for the universe $\mathcal{U}(T^b)$, both formulations have a common optimal solution.

2 Hierarchical hubbing for the generalized VPN problem

The main contribution of this note is

Theorem 3. *The optimal hierarchical hubbing solution for a generalized VPN problem can be found in polynomial time.*

Thus RND_{HH} is polynomially solvable for a large and interesting class of demand universes. A further motivation for this result is its connection with a conjecture about the polynomial solvability of the generalized VPN problem with arbitrary oblivious single path routings; this is discussed in detail in Section 3.

2.1 Preliminaries

We will use the Iverson bracket $[A]$ to denote the indicator function of a predicate A . Given two disjoint subsets A, B of vertices in some graph $G = (V, E)$, an (A, B) cut is any set $S \subseteq V$ that separates A from B . Given a tree T and edge capacities $b : E(T) \rightarrow \mathbb{R}_+$, T^b will denote the resulting edge-capacitated tree.

Definition 4. We say a tree T is a **hub tree** if its leaf set is in correspondence with the terminal set W of the instance.

Definition 5. For any hub tree T , a **T -embedding** (into G) is a map $\varphi : V(T) \cup E(T) \rightarrow V(G) \cup E(G)$ such that

- (i) $\varphi(v) \in V(G)$ for all $v \in V(T)$,
- (ii) $\varphi(i) = i$ for all $i \in W$, and
- (iii) $\varphi(vw)$ is a simple $\varphi(v)$ - $\varphi(w)$ path in G for each $vw \in E(T)$.

The restriction to simple paths in the above definition is not necessary, but will be notationally convenient; in any case, there is no advantage to using non-simple paths. But again note that for two terminals $i, j \in W$, the i - j -path induced by φ , obtained by considering the path $iv_1v_2 \cdots v_{t-1}j$ in T and concatenating the paths $\varphi(iv_1), \varphi(v_1v_2), \dots, \varphi(v_{t-1}j)$, may still be non-simple.

Definition 6. A **T -hubbing** for \mathcal{U} (into G) is a pair (φ, u) where φ is a T -embedding and u is a corresponding capacity allocation satisfying

$$u(e) \geq \sum_{f \in E(T) : e \in \varphi(f)} b(f),$$

where for each $f \in E(T)$, $b(f)$ is the maximum load on f induced by some demand in \mathcal{U} . (In reference to the earlier discussion, $b(f)$ is the capacity of the “cable” associated with edge f .)

A **hierarchical hubbing for \mathcal{U}** is simply a T -hubbing for \mathcal{U} for some choice of hub tree T .

The cost of a hierarchical hubbing solution is defined simply as the cost of the associated capacity allocation.

For a given hub tree T , it is possible that $\mathcal{U}(T^b) = \mathcal{U}(T^{b'})$ for distinct capacities b, b' . The following definition is convenient:

Definition 7. Edge capacities b for a tree T are called **defining** if for every $e \in E(T)$, there exists a $D \in \mathcal{U}(T^b)$ that saturates edge e in T^b .

Given $T^{b'}$, defining capacities b such that $\mathcal{U}(T^b) = \mathcal{U}(T^{b'})$ can easily be found, solving one maximum flow problem per edge of T .

2.2 Main theorem

We will prove (recalling the definition of $\mathcal{U}(T^b)$ from Section 1.2)

Theorem 8. *For a given hub tree T and capacities b , the optimal hierarchical hubbing for $\mathcal{U}(T^b)$ is a T -hubbing.*

We may assume in what follows that the capacities b are defining, since replacing the capacities by defining capacities can be done as a preprocessing step. Combining this theorem with the following simple but key algorithmic result about hierarchical hubbing, Theorem 3 immediately follows.

Lemma 9 ([14]). *For any hub tree T with edge capacities b , the optimal T -hubbing for $\mathcal{U}(T^b)$ can be found in polynomial time.*

The algorithm is a simple dynamic programming one which finds the optimal placement of the hubs. Note that in any optimal solution, the path $\varphi(vw)$ used to route some $vw \in E(T)$ can always be taken to be a shortest path between $\varphi(v)$ and $\varphi(w)$.

Theorem 8 is certainly very natural, and it might seem even trivial at first glance. But consider, for example, the case where T is simply a star, and hence represents a hose model universe. Then recalling Remark 1—any tree routing is a hierarchical hubbing—this theorem includes the fact that the optimal tree routing under the hose model is a hub routing, a result of Gupta et al. [10]. Further discussion on this, and connections with the “VPN Conjecture” may be found in Section 3.

We will in fact prove something stronger than Theorem 8: given any hierarchical hubbing for $\mathcal{U}(T^b)$ with capacity allocation u , there is a T -hubbing for $\mathcal{U}(T^b)$ which requires no more than capacity $u(e)$ on each edge e . We begin by proving this result for the case where the network is itself a tree, with terminals forming the leaf set of the tree. We will then observe that the result for an arbitrary network follows easily.

Tree networks Let F be an arbitrary tree, with W denoting the leaves of F . Let $\text{fund}_F(e)$ denote the set of all pairs of leaves that are separated by edge $e \in E(F)$. Notice that there is an obviously optimal oblivious routing solution for the RND problem with universe $\mathcal{U}(T^b)$ in F : use the template $\mathcal{P} = \{P_{ij} : i, j \in W\}$ where P_{ij} is the unique simple path from i to j in F . The minimal required capacity $q^*(e)$ for any edge $e \in E(F)$ is then, from (2),

$$q^*(e) = \max_{D \in \mathcal{U}(T^b)} \sum_{\{i,j\} \in \text{fund}_F(e)} D_{ij}. \quad (3)$$

Note that *any* oblivious routing template will need capacity $q^*(e)$ at least on each $e \in E(F)$, since any i - j path in F contains P_{ij} .

Observe also that this oblivious routing template is a F -hubbing template, induced by the trivial F -embedding that maps F to itself. Since each cable in this embedding is a single edge, combining this F -embedding with the capacity allocation q^* yields an F -hubbing for $\mathcal{U}(T^b)$.

The next theorem, which is the key technical theorem of this paper, shows that there is a T -hubbing for $\mathcal{U}(T^b)$ with capacity allocation q^* . In general, the T -embedding that certifies this will induce *non-simple* paths; despite this, extra capacity will not be needed.

Theorem 10. *Let F be any tree with leaf set W , and let $q^* : E(F) \rightarrow \mathbb{R}_+$ be as in (3). Then there is a T -hubbing for $\mathcal{U}(T^b)$ into F with capacity allocation q^* .*

Proof. We will need the following lemma, which follows easily from standard uncrossing techniques.

Lemma 11. *Let $W_1 \subseteq W_2 \subseteq W$, and for $i \in \{1, 2\}$, let S_i be a minimum capacity $(W_i, W \setminus W_i)$ -cut in T^b (with $W_i \subseteq S_i$). Then $S_1 \cap S_2$ is a minimum capacity $(W_1, W \setminus W_1)$ -cut, and $S_1 \cup S_2$ is a minimum capacity $(W_2, W \setminus W_2)$ -cut.*

Proof. By submodularity of the cut function,

$$b(\delta(S_1 \cap S_2)) + b(\delta(S_1 \cup S_2)) \leq b(\delta(S_1)) + b(\delta(S_2)). \quad (4)$$

But $S_1 \cap S_2$ is a $(W_1, W \setminus W_1)$ -cut, and so by the definition of S_1 , $b(\delta(S_1 \cap S_2)) \geq b(\delta(S_1))$. Moreover, $S_1 \cup S_2$ is a $(W_2, W \setminus W_2)$ -cut, and hence $b(\delta(S_1 \cup S_2)) \geq b(\delta(S_2))$. We deduce that (4) holds with equality, and hence that $S_1 \cap S_2$ is a minimum $(W_1, W \setminus W_1)$ -cut, and $S_1 \cup S_2$ a minimum $(W_2, W \setminus W_2)$ -cut. \square

Pick an arbitrary leaf $r \in W$, and call it the root. An edge e of F divides the terminal set into two, W_e and $W \setminus W_e$, where we choose W_e to *not* contain the root. Let S_e be a minimum $(W_e, W \setminus W_e)$ cut containing W_e in the tree T^b , breaking ties by choosing S_e to have minimum cardinality.

We now describe the T -embedding φ into F . We will define, for each internal node $w \in V(T) \setminus W$, an orientation \vec{F}_w of F . For each edge $e \in E(F)$, orient e away from the root if $w \in S_e$, and otherwise orient e towards the root. We then define $\varphi(w)$ to be the unique sink node of \vec{F}_w , whose existence we guarantee as follows.

Lemma 12. *There is a unique sink node in \vec{F}_w .*

Proof. We begin by showing that every node has outdegree at most 1 in \vec{F}_w . Suppose for a contradiction that some node $u \in V(F)$ has two outgoing arcs e, e' in \vec{F}_w . There are two cases to consider:

- (i) Both arcs e and e' are oriented away from r . Then $w \in S_e$ and $w \in S_{e'}$. Now $W_e \subseteq (W \setminus W_{e'})$, and hence by Lemma 11, $\hat{S}_e := S_e \cap (V(T) \setminus S_{e'})$ is a minimum $(W_e, W \setminus W_e)$ -cut in T^b . But since $w \notin \hat{S}_e$, $\hat{S}_e \subsetneq S_e$, contradicting the size minimality of S_e .
- (ii) One of e, e' is oriented towards r , say e' . Then $w \in S_e$ and $w \notin S_{e'}$. Then simply note that $W_e \subseteq W_{e'}$, and hence (by Lemma 11) $S_e \cap S_{e'}$ is a minimum $(W_e, W \setminus W_e)$ -cut in T^b . Again since $w \notin S_e \cap S_{e'}$, this contradicts the size minimality of S_e .

By starting at an arbitrary node, and following the unique outgoing arc until we reach a node with no outgoing arcs (this may be a leaf, or not), the existence of some sink v follows. To see that this sink is unique, observe that on any path in \vec{F}_w terminating at v , all arcs must be oriented towards v by the condition on the outdegree, and hence none of the nodes on the path aside from v can possibly be a sink. \square

We complete the definition of the T -embedding φ in the obvious way, by taking $\varphi(vw)$ to be the unique simple path between $\varphi(v)$ and $\varphi(w)$ in F , for each $vw \in E(T)$.

Now let us consider $q^*(e)$ for some edge $e \in E(F)$. Rewriting (3), we have

$$q^*(e) = \max_{D \in \mathcal{U}(T^b)} \sum_{i \in W_e, j \notin W_e} D_{ij}. \quad (5)$$

The right hand side of (5) can be seen as a maximum flow problem; send as much flow as possible in T^b from W_e to $W \setminus W_e$. Invoking the max-flow min-cut theorem, we obtain

$$q^*(e) = b(\delta(S_e)). \quad (6)$$

We show now that the T -embedding φ along with the capacity allocation q^* defines a valid hierarchical hubbing solution. Recall that $[\cdot]$ denotes the Iverson bracket. Consider any edge e in F . The capacity required on edge e by the hierarchical hubbing solution induced by φ is

$$\begin{aligned} & \sum_{f=uv \in E(T)} [\varphi(u) \text{ and } \varphi(v) \text{ are separated by } e] b(f) \\ &= \sum_{f=uv \in E(T)} [\vec{F}_u \text{ and } \vec{F}_v \text{ orient } e \text{ in opposite directions}] b(f) \\ &= \sum_{f=uv \in E(T)} [\text{exactly one of } u \text{ and } v \text{ is in } S_e] b(f) \\ &= b(\delta(S_e)). \end{aligned}$$

Combined with (6), this completes the proof. \square

General networks We now show how this result for tree networks can be leveraged to demonstrate the general case.

Let (φ, u) be any F -hubbing for $\mathcal{U}(T^b)$ into G . By Theorem 10, there exists a T -hubbing (η, q^*) for $\mathcal{U}(T^b)$ into F , where q^* is the minimal possible capacity allocation discussed earlier. We will essentially compose this hierarchical hubbing with φ to obtain a T -hubbing (ρ, u) into G . We define the T -embedding ρ as follows. Let $\rho(v) = \varphi(\eta(v))$ for all $v \in V(T)$. For any edge $vw \in E(T)$, consider the path $\eta(vw)$ in F , and write it in terms of its edges: $\eta(vw) = e_1 e_2 \dots e_t$. Take $\rho(vw)$ to be any simple $\rho(v)$ - $\rho(w)$ path in G contained in the concatenation of the paths $\varphi(e_1), \varphi(e_2), \dots, \varphi(e_t)$. Clearly ρ does define a T -embedding into G .

We have that for any $e \in E(G)$,

$$u(e) \geq \sum_{f \in E(F)} [e \in \varphi(f)] q^*(f).$$

This follows from the definition of a F -hubbing for $\mathcal{U}(T^b)$, combined with (3). Now since (η, q^*) is a T -hubbing for $\mathcal{U}(T^b)$ into F , we have that

$$q^*(f) \geq \sum_{f' \in E(T)} [f' \in \eta(f)] b(f').$$

Hence

$$\begin{aligned}
u(e) &\geq \sum_{f' \in E(T)} \sum_{f \in E(F)} [e \in \varphi(f)] \cdot [f' \in \eta(f)] b(f') \\
&\geq \sum_{f' \in E(T)} [e \in \rho(f')] b(f').
\end{aligned}$$

So (ρ, u) is indeed a T -hubbing for $\mathcal{U}(T^b)$ into G . The proof of Theorem 8, and hence Theorem 3, is complete.

3 Connection to the VPN Conjecture

A well-known conjecture in the area was the *VPN Conjecture*. This was resolved by Goyal et al. [9], who proved

Theorem 13 ([9]). *There is an optimal solution to the RND problem with oblivious routing for the hose model whose support is a tree.*

This result leads to a polynomial time algorithm to solve the RND problem for the hose model exactly: indeed, Gupta et al. [10] had previously provided a polynomial time algorithm that computes the optimal tree solution. The algorithm simply finds a hub routing solution. In other words, for each node v , compute the sum (weighted by the hose capacities b_i) of the lengths of the shortest paths from each terminal to v ; then choosing v^* that minimizes this quantity, allocate b_i units of capacity along the i - v^* shortest path for each $i \in W$, additively¹. The routing is simply a hub routing centered at v^* : for each pair i, j of terminals, the path from i to j is obtained by appending the shortest path from i to v^* to that from v^* to j . This is of course a $R^{b'}$ -hubbing solution, where R is a star with leaf set W , and the capacity of edge ir (with r being the internal node of the star) is $b'(ir) = b_i$.

The result of Gupta et al. [10] can be seen as a precursor to, and evidence for, the VPN Conjecture. Similarly, the result of this paper, which is a generalization of the result of Gupta et al., can be seen as evidence for the following conjecture proposed in [14].

Generalized VPN Conjecture. *There is an optimal solution to the RND problem with universe $\mathcal{U}(T^b)$ that is a T -hubbing solution.*

¹A technicality: the marginals should be “defining” in the sense of Definition 7; here, this means that $b_i \leq \frac{1}{2} \sum_j b_j$ for all $i \in W$.

Theorem 3 shows that the following version of the conjecture, while seemingly weaker, is equivalent:

Generalized VPN Conjecture (II). *There is an optimal solution to the RND problem with universe $\mathcal{U}(T^b)$ that is a hierarchical hubbing solution.*

It is known that the optimal T -hubbing solution for $\mathcal{U}(T^b)$ is always within a constant factor of the optimal oblivious routing solution [14]. This in particular implies that the RND problem for the generalized hose model is constant approximable. A positive resolution to the Generalized VPN Conjecture would imply that the RND problem for $\mathcal{U}(T^b)$ can be solved optimally in polynomial time.

Acknowledgements The author is very grateful to Alexandre Fréchette and Bruce Shepherd, both for asking the question that led to this note, and for their help in the exposition of the main proof.

References

- [1] W. Ben-Ameur and H. Kerivin. New economical virtual private networks. *Communications of the ACM*, 46(6):69–73, June 2003.
- [2] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, August 2009.
- [3] G. Calinescu, H. Karloff, and Y. Rabani. Approximation algorithms for the 0-extension problem. *SIAM Journal on Computing*, 34(2):358–372, 2005.
- [4] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. van der Merwe. A flexible model for resource management in virtual private networks. In *Proceedings of SIGCOMM*, pages 95–108, New York, NY, USA, 1999. ACM.
- [5] F. Eisenbrand and E. Happ. Provisioning a virtual private network under the presence of non-communicating groups. In *Proceedings of the 6th International Conference on Algorithms and Complexity*, volume 3998 of *Lecture Notes in Computer Science*, pages 105–114. Springer, 2006.

- [6] J. A. Fingerhut, S. Suri, and J. S. Turner. Designing least-cost non-blocking broadband networks. *Journal of Algorithms*, 24(2):287–309, August 1997.
- [7] A. Fréchette. Hub routing for the robust network design problem. Master’s thesis, McGill University, 2012.
- [8] A. Fréchette, F. B. Shepherd, M. K. Thottan, and P. J. Winzer. Shortest path versus multi-hub routing in networks with uncertain demand. In *Proceedings of IEEE INFOCOM*, pages 710–718, 2013.
- [9] N. Goyal, N. Olver, and F. B. Shepherd. The VPN conjecture is true. *J. ACM*, 60(3):17:1–17:17, June 2013.
- [10] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener. Provisioning a virtual private network: a network design problem for multi-commodity flow. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 389–398, New York, NY, USA, 2001. ACM Press.
- [11] A. Karzanov. Minimum 0-extensions of graph metrics. *European Journal of Combinatorics*, 19(1):71–101, January 1998.
- [12] A. Kumar, R. Rastogi, A. Silberschatz, and B. Yener. Algorithms for provisioning virtual private networks in the hose model. *IEEE/ACM Transactions on Networking*, 10(4):565–578, August 2002.
- [13] N. Olver. *Robust Network Design*. PhD thesis, McGill University, 2010.
- [14] N. Olver and F. B. Shepherd. Approximability of robust network design. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1097–1105, 2010.
- [15] F. B. Shepherd and P. J. Winzer. Selective randomized load balancing and mesh networks with changing demands. *Journal of Optical Networking*, 5(5):320–339, May 2006.